Bank of Canada Banque du Canada

Working Paper 2001-22 / Document de travail 2001-22

On Inflation and the Persistence of Shocks to Output

by

Maral Kichian and Richard Luger

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2001-22

December 2001

On Inflation and the Persistence of Shocks to Output

by

Maral Kichian and Richard Luger

Research Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
mkichian@bankofcanada.ca
rluger@bankofcanada.ca

The views expressed in this paper are those of the authors.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgement iv
Abstract/Résumé v

1. Introduction 1

2. The Threshold ARMA Model 3

   2.1 Diagnostic checks 5

3. Testing for Threshold Effects 8

4. Impulse Responses 10

   4.1 The role of inflation 11

   4.2 Asymmetry of output shocks 13

5. Conclusion 14

Bibliography 16

Appendix 18

Acknowledgement

The authors would like to thank Alain Guay for helpful comments.

Abstract

This paper empirically investigates the possibility that the effects of shocks to output depend on the level of inflation. The analysis extends Elwood's (1998) framework by incorporating in the model an inflation-threshold process that can potentially influence the stochastic properties of output. The value of this threshold parameter, if it exists, is considered to be unknown and is estimated in the model. The results indicate that shocks to output indeed have asymmetric effects, depending on the level of inflation: negative shocks are more detrimental when inflation is high, and positive shocks are more persistent when inflation is low.

JEL classification: E31, E32, E52, E58
Bank classification: Econometric and statistical methods; Inflation: costs and benefits

Résumé

Les auteurs cherchent à établir empiriquement si l'incidence des chocs de production varie selon le niveau de l'inflation. Ils prolongent l'analyse d'Elwood (1998) en intégrant au modèle un seuil d'inflation susceptible d'influer sur les propriétés stochastiques de la production. La valeur de ce seuil étant inconnue, elle est estimée dans le cadre du modèle. D'après les résultats obtenus, les chocs de production ont effectivement des effets asymétriques, qui dépendent du niveau de l'inflation : les chocs négatifs sont plus dommageables lorsque l'inflation est élevée, et les chocs positifs sont plus persistants quand l'inflation est faible.

Classification JEL : E31, E32, E52, E58
Classification de la Banque : Méthodes économétriques et statistiques; Inflation : coûts et avantages

102-01558/

1. Introduction

This paper empirically examines whether the level of inflation matters for the persistence of output growth. The idea that inflation could have such threshold effects is worth investigating because some authors have suggested that a low-inflation environment has been instrumental in generating the unprecedented and sustained output-growth rates recently experienced by some countries. For instance, Taylor (1998) notes that the United States experienced its two longest post-war expansions after 1983, and he suggests that monetary policy was the main cause: "by keeping the inflation low and stable . . . the Fed [Federal Reserve] has succeeded in stabilizing the economy and making recessions less frequent, smaller, and shorter."

Studies looking for other types of non-linearities in the output-growth process include that by Beaudry and Koop (1993). They show that positive shocks generate substantially different output dynamics than negative shocks when an index variable that captures the depth of recessions is included in a standard autoregressive moving average (ARMA) model for output. They conclude that, if shocks of opposing signs do indeed have asymmetric effects, imposing symmetry will bias the estimates of the persistence parameters of output.

In this paper we thus account for the possibility that both inflation and shocks influence output-growth behaviour, and that this impact can vary with the level of inflation and the sign of output-growth shocks. At this stage, our aim is simply to document whether such non-linearities exist. We do not examine the important issue of why they might exist. The corresponding theoretical analysis is left for future research.

The model that we propose is a generalization of Elwood's (1998) unobserved-components threshold framework which, like Beaudry and Koop, was initially used to examine whether disturbances have asymmetric effects on output growth. A main advantage of Elwood's methodology over Beaudry and Koop's is that,

instead of using a possibly imperfect proxy, shocks are allowed to directly influence output-growth behaviour.[1] More precisely, the disturbance term is treated as an unobserved component and its estimated sign determines the regime within which output growth evolves. On the other hand, Beaudry and Koop use ARMA specifications, which are more flexible than Elwood's low-order AR or MA models. A more general framework, combining elements from both studies, could prove to be a useful modelling alternative.

We retain the main Elwood structure but extend it to an ARMA setting. In addition, to integrate the possible role of the inflation environment on output growth, we further generalize the model by allowing for multiple threshold effects. Consequently, the parameters of output growth are permitted to change, depending on (i) whether disturbances are positive or negative, and (ii) whether inflation is above or below some threshold level. We then test for these distinct effects using Canadian data, since Canada has had an announced low-inflation policy from the early 1990s. Maximum-likelihood estimation is used, and hypotheses are tested using Hansen's (1996) bootstrap test procedure for when a nuisance parameter is present only under the alternative.

Our results concur with the conclusions of Beaudry and Koop that shocks indeed have asymmetric effects on output. However, we show that the inflation environment at the time of the shock plays a crucial role in determining which of the shocks displays the greater persistence. Thus, under low inflation, a positive shock is found to be more persistent than a negative shock of the same size. The reverse is true when inflation is above its threshold value. Therefore, low inflation is associated with healthier output-growth dynamics than high inflation. The inflation-threshold level of the model was estimated to be 4.4 per cent.

Section 2 presents our generalized multiple-threshold framework and reports

[1] This might be partly the reason why, for the same data, Elwood did not find significantly different effects.

the maximum-likelihood estimation results. Section 3 tests for the threshold effect using the likelihood ratio statistic. Since the threshold parameter is not identified under the null hypothesis, the asymptotic distribution of the test statistic is simulated using Hansen's (1996) procedure. The test results indicate that the threshold effect is indeed significant. In section 4, impulse-response functions show the effects that one-time positive and negative shocks have on output growth and level. Section 5 concludes.

2. The Threshold ARMA Model

The approach that we adopt to investigate the possibility of a threshold effect in the inflation-output relation is based on the class of threshold autoregressive models introduced by Tong (1978). In such models, changes in parameter values are endogenously generated by a fixed lag of the *observed* series. See Tong and Lim (1980) and Tong (1983, 1990) for more details.

Elwood (1998) proposed an extension to threshold ARMA models where the sign of a fixed lag of the *unobserved* shocks determines how the output-growth parameters change. His unobserved-components methodology thus provides a framework for detecting asymmetries in the persistence of shocks to output. To investigate the additional threshold effects of inflation, we consider the following four-regime threshold ARMA model:

$$\Delta y_t = \mu + \phi_1^p(\Delta y_{t-1} - \mu) + \varepsilon_t + \theta_1^p \varepsilon_{t-1}, \quad \varepsilon_{t-1} \geq 0, \ \pi_{t-1} \geq d, \tag{1}$$

$$\Delta y_t = \mu + \phi_1^n(\Delta y_{t-1} - \mu) + \varepsilon_t + \theta_1^n \varepsilon_{t-1}, \quad \varepsilon_{t-1} < 0, \ \pi_{t-1} \geq d, \tag{2}$$

$$\Delta y_t = \mu + \phi_2^p(\Delta y_{t-1} - \mu) + \varepsilon_t + \theta_2^p \varepsilon_{t-1}, \quad \varepsilon_{t-1} \geq 0, \ \pi_{t-1} < d, \tag{3}$$

$$\Delta y_t = \mu + \phi_2^n(\Delta y_{t-1} - \mu) + \varepsilon_t + \theta_2^n \varepsilon_{t-1}, \quad \varepsilon_{t-1} < 0, \ \pi_{t-1} < d, \tag{4}$$

where Δy_t is output growth, π_t is the rate of inflation, and $\varepsilon_t \sim i.i.d.\ N(0, \sigma^2)$. Given this specification, output growth Δy_{t+1} depends on the sign of the time,

t, shock, ε_t, and the level of inflation, π_t. Therefore, the persistence of positive and negative shocks may differ depending on the values of the parameters ϕ_j^i, θ_j^i, $j = 1, 2$, $i = p, n$. For example, if $\phi_2^p + \theta_2^p > \phi_2^n + \theta_2^n$ and $\phi_1^p + \theta_1^p < \phi_1^n + \theta_1^n$, then positive shocks have more persistent effects on output growth than do negative shocks when inflation is low, and vice versa when inflation is high.

For a given value of the inflation threshold level, d, the remaining model parameters can be estimated using the modified Kalman filter proposed by Elwood (1998). The appendix provides the state-space representation of the threshold ARMA(1,1) model described above, and the estimation methodology. Denote the maximized-likelihood function for a given value of d by $\hat{L}_1(d)$. The estimate for d can then be defined as

$$\hat{d} = \arg\max_{d \in \mathcal{D}} \hat{L}_1(d), \tag{5}$$

where $\mathcal{D}$ is the set of admissible values for d. The range of admissible values for the inflation-threshold parameter is defined as the observed range of inflation levels, with 15 per cent trimmed at both ends.[2]

In principle, higher-order threshold ARMA models could be considered by extending the estimation methodology. However, the number of parameters grows exponentially with the number of regimes. Because our goal is to simply investigate whether the data support the presence of threshold effects, we limit our analysis to the first-order case. Nevertheless, despite its apparent simplicity, the threshold ARMA(1,1) is a parsimonious representation of a potentially highly asymmetric time series. The proposed model can be decomposed into four unobserved components, each receiving all the shocks that are specific to their sign and inflation regime:

$$\Delta y_t = \mu + \Delta y_{t1}^p + \Delta y_{t1}^n + \Delta y_{t2}^p + \Delta y_{t2}^n + \varepsilon_t, \tag{6}$$

where Δy_{tj}^i are infinite moving-average processes for $j = 1, 2$, and $i = p, n$. Thus, for example, the first unobserved component is the infinite sum of all past positive

[2]The trimming rule follows Andrews (1993) and Hansen (1996).

shocks that occurred while inflation was high:

$$\Delta y_{t1}^{p} = \frac{\theta_1^p B}{1 - \phi_1^p B}\varepsilon_t, \text{ for all } \varepsilon_{t-k} \geq 0,\ \pi_{t-k} \geq d, \tag{7}$$

defined for $k \geq 1$ and where B is the lag operator such that $B^k z_t = z_{t-k}$. The other components in (6) are defined in a similar fashion.[3]

The proposed model in equations (1) through (4) was estimated using Canadian data on real GDP and on the associated implicit prices, over the period 1965Q1 to 2000Q3.[4] Specifically, the growth rate of output is the annualized log difference of seasonally adjusted real GDP, while inflation is the annualized log difference of the GDP deflator. Trimming 15 per cent of the highest and lowest values of the inflation series yielded the interval [1.5,7.5] for $\mathcal{D}$, over which we defined a grid of 60 possible values that vary by increments of 0.1. The value of the inflation-threshold parameter estimated by the grid-search method resulted in $\hat{d} = 4.40$, which is statistically significant, as we will see in section 3. Table 1 summarizes the estimation results for the remainder of the model parameters. Figure 1 shows the estimated inflation-threshold level against the output and inflation series.

2.1 Diagnostic checks

The quasi-maximum-likelihood parameter estimates presented in Table 1 were used in the modified Kalman filter to obtain residuals, e_t (see equation (31) in the appendix). Using this series, a number of statistics were computed to test the model's specifications.

[3]Seen this way, the proposed model is in fact an extension of the asymmetric moving-average model developed by Wecker (1981), as were the models proposed by Elwood (1998).

[4]The sample starting date was chosen to strike a balance between having adequate economic information and avoiding possible structural changes early on in the data. That is, we include the essence of the economic environment prior to the high-inflation era of the 1970s, but do not start earlier than 1965, given the robustness concerns in an empirical investigation of this nature.

Table 1
Estimation Results

μ	ϕ_1^p	ϕ_2^p	θ_1^p	θ_2^p	ϕ_1^n	ϕ_2^n	θ_1^n	θ_2^n	σ
4.13	0.86	-0.56	-0.99	0.99	-0.31	0.96	0.98	-0.99	3.31
(0.04)	(0.22)	(0.005)	(0.06)	(0.09)	(0.006)	(0.71)	(0.07)	(0.06)	(0.03)

Notes: Superscript p refers to coefficients of positive shocks, superscript n refers to coefficients of negative shocks, subscript 1 refers to a high-inflation regime, while subscript 2 indicates coefficients estimated for the low-inflation regime. The numbers in parentheses are the asymptotic standard errors derived from the elements along the main diagonal of the inverse of the information matrix.

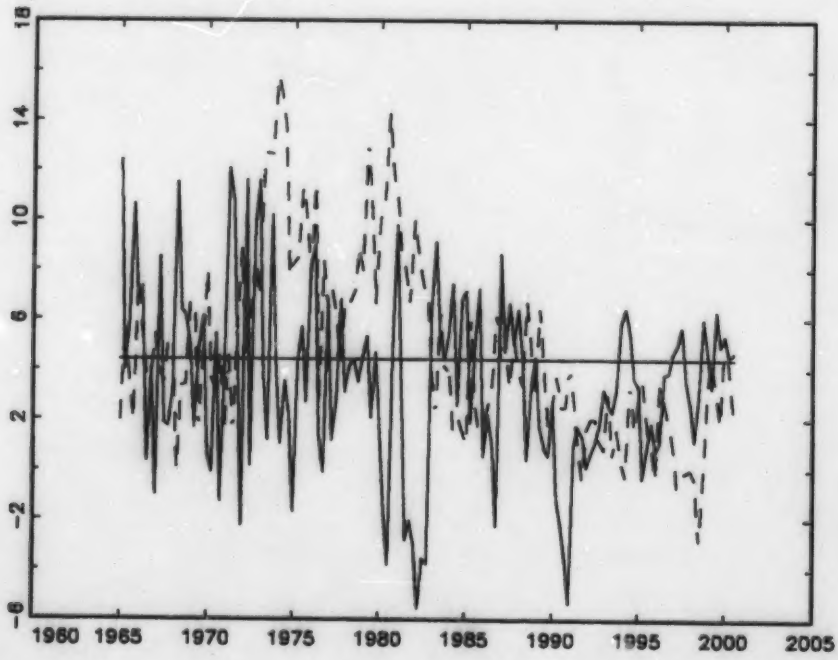

Figure 1: Growth rates of Canadian GDP (solid line) and the GDP deflator (dashed line). The horizontal line represents the estimated inflation-threshold level.

The statistic z_N, which asymptotically is χ^2_2, is the Bowman and Shenton (1975) test for residual non-normality computed as

$$z_N = (T/6)\left(\hat{\sigma}^{-3}\sum_{t=1}^{T} e_t^3/T\right)^2 + (T/24)\left(\hat{\sigma}^{-4}\sum_{t=1}^{T} e_t^4/T - 3\right)^2, \qquad (8)$$

where $\hat{\sigma}^2 = \sum_{t=1}^{T} e_t^2/T$. The computed value of 6.50 has an associated p-value of 0.041, which is marginally significant at the conventional 5 per cent level. The statistic z_H is a Goldfeld and Quandt (1965) type heteroscedasticity test computed as

$$z_H = \sum_{t=T-m+1}^{T} e_t^2 \Big/ \sum_{t=1}^{m} e_t^2, \qquad (9)$$

with $m = T/3$, and where mz_H is asymptotically χ^2_m. Finally, the statistics $z_{Q(p)}$, for $p = 1, ..., 6$, are Ljung and Box (1979) autocorrelation tests computed as

$$z_{Q(i)} = T(T+2)\sum_{h=1}^{p}(T-h)^{-1}\hat{\rho}(h)^2, \qquad (10)$$

and distributed asymptotically as χ^2_p, where $\hat{\rho}(h) = \sum_{t=h+1}^{T} e_t e_{t-h} / \sum_{t=1}^{T} e_t^2$ is the sample autocorrelation. See Harvey (1990, section 5.2) for more on these tests. Table 2 presents the results of these diagnostic checks. Besides the slight departure from normality, the residuals seem to support the model specification.

Table 2
Diagnostic Tests

z_N	mz_H	$z_{Q(1)}$	$z_{Q(2)}$	$z_{Q(3)}$	$z_{Q(4)}$	$z_{Q(5)}$	$z_{Q(6)}$
6.50*	10.78	0.45	2.87	5.90	8.69	8.82	10.52

Notes: The statistic z_N is a test for residual non-normality, which asymptotically is χ^2_2. The star indicates statistical significance at the 5 per cent level. The statistic z_H is a heteroscedasticity test for which mz_H is asymptotically χ^2_m, with $m = T/3$. Finally, $z_{Q(p)}$ are autocorrelation tests based on p lags, for $p = 1, ..., 6$, that asymptotically are χ^2_p.

3. Testing for Threshold Effects

There are two hypotheses of interest in this model. The first is that there is no inflation-threshold effect; i.e., the level of inflation has no effect on output growth. In the notation for the model, this hypothesis is represented as

$$H_{01} : \phi_1^i = \phi_2^i \text{ and } \theta_1^i = \theta_2^i, \quad \text{for } i = p, n. \tag{11}$$

The second hypothesis is that shocks of different signs have similar effects on output. That is,

$$H_{02} : \phi_j^p = \phi_j^n \text{ and } \theta_j^p = \theta_j^n, \quad \text{for } j = 1, 2. \tag{12}$$

Statistical testing of H_{02} can be performed on the basis of standard asymptotic theory. We thus compute the usual likelihood ratio statistic defined as $LR_{02}(d) = 2\log(\hat{L}_1(d)/\hat{L}_{02}(d))$, where we recall that $\hat{L}_1(d)$ is the maximized value of the unconstrained likelihood function, while $\hat{L}_{02}(d)$ is that under H_{02}. Under H_{02}, the likelihood ratio LR_{02} is asymptotically distributed as $\chi^2(r)$, with degrees of freedom $r = 4$. We find that $LR_{02}(\hat{d}) = 10.26$, which has a p-value of 0.0375, indicating that shocks of different signs have significantly different effects on output.

Statistical testing of H_{01} is not as straightforward. Under the null hypothesis of a no-inflation-threshold effect, the threshold parameter d is not identified. In such a case, standard asymptotic inference is invalid, since the information matrix is singular under the null hypothesis. To account for the fact that some parameters are present only under the alternative, we use Hansen's (1996) bootstrap test procedure. It allows us to simulate the limiting distribution of the supremum likelihood ratio that results from a maximization over the space of the threshold parameter. The bootstrap critical values were based on 1,000 replications of the simulation procedure.

Figure 2 plots the likelihood ratio statistic $LR_{01}(d) = 2\log(\hat{L}_1(d)/\hat{L}_{01})$ for the 60 values of d, where $\hat{L}_{01}$ is the maximized value of the likelihood function under

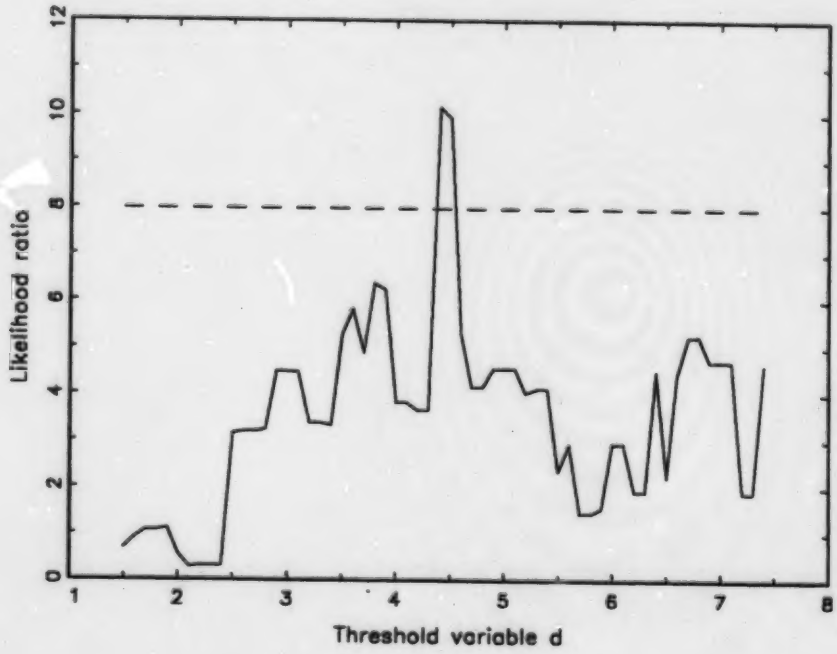

Figure 2: The solid line corresponds to the value of the likelihood function for a given value of the inflation-threshold variable, d. The dashed line is the bootstrap 5 per cent critical value. Values for which the likelihood is above the dashed line yield the 95 per cent confidence interval for d.

H_{01}. Obviously, $LR_{01}(d)$ reaches its maximum value at the threshold estimate $\hat{d} = 4.40$. The dotted line in the graph represents the bootstrap critical value at the 5 per cent level. Thus, values for which the likelihood is above the dashed line yield the 95 per cent confidence interval for d. Figure 2 shows that the confidence interval is quite tight around the threshold estimate. It appears that output-growth behaviour does indeed depend upon the level of inflation.

4. Impulse Responses

In this section we describe impulse responses for the growth rate and the level of output. Impulse responses can be computed by taking the difference between a shocked and a base case. Since our model is non-linear, these functions depend both on the lagged values of output before the shock and on the size of the imputed shock. Following Beaudry and Koop (1993), we construct unconditional impulse responses, in the sense that we compare the after-shock effect with a base case where output growth equals its mean and where all past disturbances are zero. In addition, we normalize the shocks to have a unit variance, so that a shock of one unit is equivalent to a shock of one standard deviation. Thus, the impulse response function of output growth to a shock v at time $t = 0$ for, say, the high-inflation regime, τ periods ahead, is given by

$$RF(\Delta y_\tau; v) = E[\Delta y_\tau | \Delta y_0 = \mu + v, \varepsilon_0 = v, \pi_\tau \geq d] - \mu, \tag{13}$$

defined for $\tau \geq 1$, and where we set $v = \sigma = 1$. The response function of the level of output is then obtained simply as

$$RF(y_\tau; v) = RF(\Delta y_\tau; v) + RF(y_{\tau-1}; v), \tag{14}$$

with $RF(y_0; v) = v$. Therefore, conditional on a given inflation regime, the impulse responses are independent of the history of the time series, as in the case of a

standard linear model. However, unlike the standard case, the impulse responses depend on the sign of the imputed shock.[5]

Figures 3 and 4 show the impulse responses of the growth rate and the level of output, respectively. The solid lines represent the responses to shocks when inflation is below its threshold level, and the dashed lines represent the corresponding responses when inflation is above its threshold level.

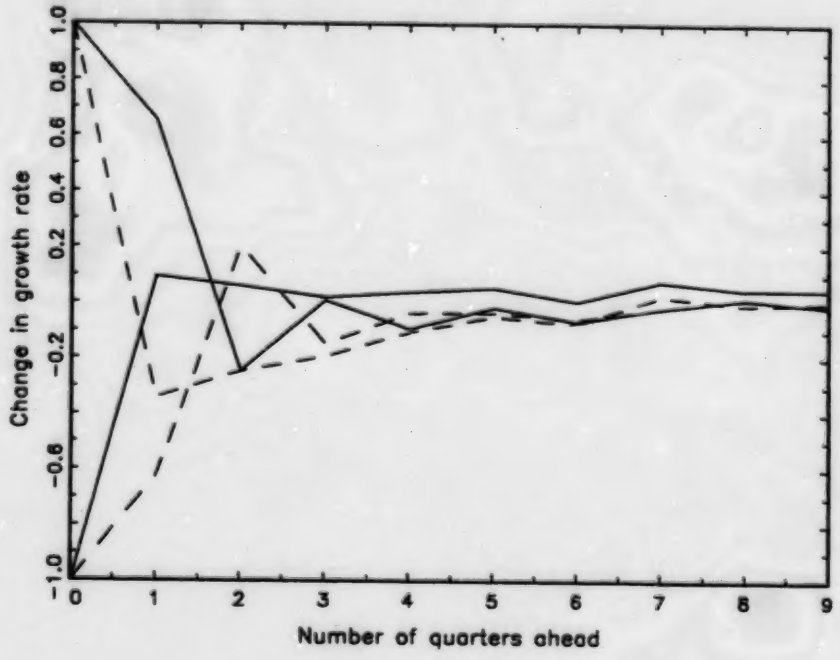

Figure 3: Impulse response of the growth rate of output. The solid lines are associated with shocks in the low-inflation regime, while the dashed lines are associated with shocks in the high-inflation regime.

4.1 The role of inflation

The impact of a unit positive shock on output growth is twice as persistent in a low-inflation regime as it is in a high-inflation regime: the effect on output

[5]See Potter (1995) for more on non-linear impulse response functions.

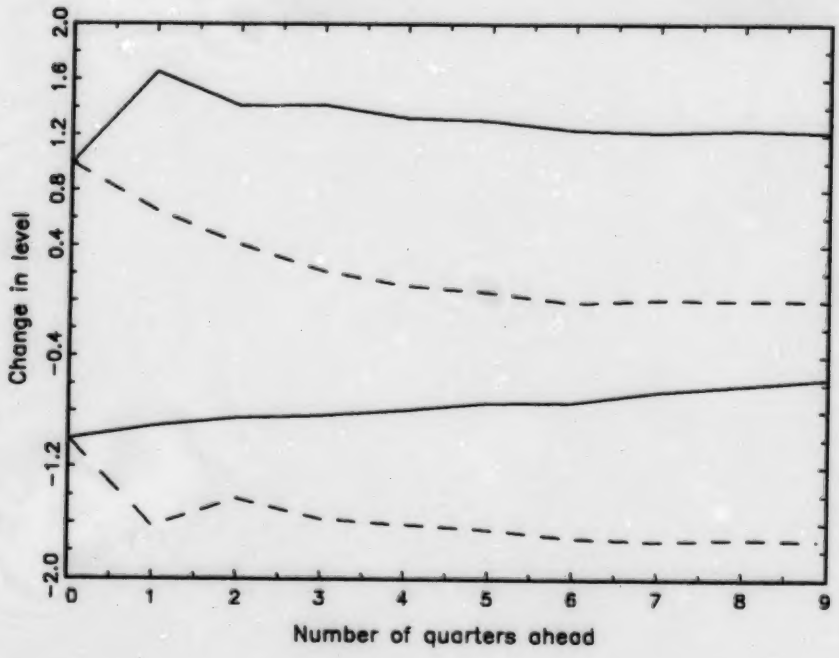

Figure 4: Impulse response of the level of output. The solid lines are associated with shocks in the low-inflation regime, while the dashed lines are associated with shocks in the high-inflation regime.

growth is felt over two quarters instead of just one. On the other hand, a negative shock causes the growth rate to rebound twice as quickly to its mean level when inflation is in the low- rather than high-inflation regime. Thus, a negative shock is less persistent in a time of low inflation than it is in a time of high inflation.

For the impulse responses on the output level, we find that a positive shock of one standard deviation, in a low-inflation regime, causes output to increase substantially and to settle 2 years later at a level that is 20 per cent higher than the amount of the shock. In contrast, the same positive shock under a high-inflation regime will dissipate within $1\frac{1}{2}$ years. Clearly, a positive shock has a much more beneficial effect on the economy when it arrives in a time of low inflation. The effect of a negative one-unit shock is equally telling: under high inflation, it causes output to decline considerably, such that 2 years later output will have diminished by $1\frac{1}{2}$ times the amount of the shock. In contrast, with low inflation, the effect of the shock decreases, albeit at a slow rate.

We can therefore conclude that a low-inflation regime is clearly more desirable. Of course, the low level of inflation may not in itself be the cause of the good times described above; rather, the cause may be that the variable is able to capture an underlying set of structural non-linear conditions that are favourable to the economy.

4.2 Asymmetry of output shocks

For a given inflation level, we now compare the response of the GDP to positive and negative shocks of equal size. Consider first a regime where inflation is above its threshold value (the dashed lines in Figure 4). It is easy to see that a positive standard-deviation shock to output is less persistent than an equivalent negative shock. In fact, while the former has already dissipated $1\frac{1}{2}$ years after the initial impact, the negative shock causes output to fall by an extra half standard deviation immediately, and commits GDP to that low level well beyond 2 years. For the reaction of output to these shocks in a low-inflation regime (the solid lines in Figure 4), again there is asymmetry. Two years after a one-unit positive shock, output rises by an additional 20 to 25 per cent of the amount of the shock. In contrast, the negative shock is not exacerbated, but slowly starts to dissipate over time.

It can therefore be concluded that, for a given inflation regime, shocks to output have asymmetric effects. Furthermore, this asymmetry is more pronounced in a high-inflation regime. Interestingly, while a negative shock has a larger size-effect under high inflation than a positive shock has under low inflation, both shocks display similar dynamics over time.

5. Conclusion

A number of researchers have suggested that the sustained and strong output-growth levels observed in numerous countries over the past decade are mainly attributable to the existence of a low and stable inflation regime in those countries. The main purpose of this paper was to examine this question empirically in the case of Canada.

Our methodology consisted of extending Elwood's (1998) unobserved-components framework. Thus, in addition to having the sign of lagged shocks determine how output-growth parameters change, we allowed for the possibility that a second threshold effect on output growth arises from the inflation level. Output growth was permitted to evolve according to four possible regimes, which depended on the sign of the lagged output-growth shock and on whether inflation was above or below some critical level. We then estimated an ARMA(1,1) specification with the above assumptions using Canadian data on GDP growth and inflation. The estimation methodology used the Kalman filter; maximum-likelihood estimates were obtained for the persistence parameters, the mean and variance of the output-growth series, and for the inflation-threshold level.

Results from a standard likelihood ratio test confirmed that positive and negative shocks have significantly different effects on output, similar to the conclusion reached by Beaudry and Koop (1993). However, to test the hypothesis of no-inflation-threshold, it was necessary to simulate the distribution of the likelihood ratio statistic (following the bootstrap procedure developed by Hansen 1996), which is non-standard, because the threshold is not identified under the null hypothesis. The results from this test showed that the inflation-threshold effect was also significant.

The above findings were summarized by calculating of impulse-response functions for the different inflation regimes. These showed that shocks to output, whether negative or positive, could be long-lived or temporary, depending on the

inflation regime. In particular, we found that a positive shock had a permanent effect on output in a low-inflation regime, while a negative shock was highly persistent in a high-inflation regime. Similarly, a positive output shock had a temporary effect when inflation was above its threshold value, whereas a negative shock had more temporary effects when inflation was low.

These results might shed some light on the very different conclusions that Campbell and Mankiw (1987) and Clark (1987) reached on the behaviour of U.S. GNP. Using linear autoregressive integrated moving average (ARIMA) models, and without distinguishing between positive or negative shocks, Campbell and Mankiw found that a 1 per cent innovation to current output changed the long-run forecast of this series by more than 1 per cent. However, using a more restricted ARIMA model, Clark found that these shocks were of a more temporary nature. In fact, if the U.S. series has the same type of non-linearities as the ones we have explored, then both of these studies could be capturing only a part of the behaviour of output. We leave this question for future research.

Bibliography

Andrews, D.W.K. 1993. "Tests for Parameter Instability and Structural Change with Unknown Change Point." *Econometrica* 55: 1465-71.

Beaudry, P. and G. Koop. 1993. "Do Recessions Permanently Change Output?' *Journal of Monetary Economics* 31: 149-63.

Bowman, K.O. and L.R. Shenton. 1975. "Omnibus Test Contours for Departures from Normality Based on $\sqrt{b_1}$ and b_2." *Biometrika* 62: 243-50.

Campbell, J. and G. Mankiw. 1987. "Are Output Fluctuations Transitory?" *Quarterly Journal of Economics* 102: 857-80.

Clark, P. 1987. "The Cyclical Component of US Economic Activity." *Quarterly Journal of Economics* 102: 798-814.

Elwood, S.K. 1998. "Is the Persistence of Shocks to Output Symmetric?" *Journal of Monetary Economics* 41: 411-26.

Goldfeld, S. and R. Quandt. 1965. "Some Tests for Homoscedasticity." *Journal of the American Statistical Association* 60: 539-47.

Hamilton, J.D. 1994. "State-Space Models." In *Handbook of Econometrics*, volume 4, edited by R. Engle and D. McFadden. Amsterdam: North-Holland.

Hansen, B. 1996. "Inference When a Nuisance Parameter is not Identified Under the Null Hypothesis." *Econometrica* 64(2): 413-30.

Harvey, A.C. 1990. *The Econometric Analysis of Time Series.* Cambridge: The MIT Press.

Ljung, T. and G. Box. 1979. "The Likelihood Function for a Stationary Autoregressive Moving Average Process." *Biometrika* 66: 265-70.

Potter, S. 1995. "A Nonlinear Approach to US GNP." *Journal of Applied Econometrics* 10: 109-25.

Taylor, J. 1998. "Monetary Policy and the Long Boom." *Review of Federal Reserve Bank of St. Louis* (November/December): 3-11.

Tong, H. 1978. "On a Threshold Model." In *On a Threshold Model, Pattern Recognition and Signal Processing*, edited by C. H. Chen. Amsterdam: Sijthoff and Noordhoff.

Tong, H. and K. Lim. 1980. "Threshold Autoregression, Limit Cycles and Cyclical Data." *Journal of the Royal Statistical Society* Series B 42: 245-92.

Tong, H. 1983. *Threshold Models in Nonlinear Time Series Analysis.* New York: Springer.

Tong, H. 1990. *Non-linear Time Series: A Dynamical Systems Approach.* Oxford: Oxford University Press.

Wecker, W.E. 1981. "Asymmetric Time Series." *Journal of the American Statistical Association* 76: 16-21.

Appendix

This appendix reviews Elwood's (1998) modified Kalman filter as used for the estimation of the threshold ARMA(1,1). The state-space representation comprises the observation or measurement equation,

$$\Delta y_t = \mu + [\,1 \quad 0\,] \begin{bmatrix} S_t \\ \varepsilon_t \end{bmatrix} \tag{15}$$

with the state vector, $[S_t\ \varepsilon_t]'$, governed by the transition equations

$$\begin{bmatrix} S_t \\ \varepsilon_t \end{bmatrix} = \begin{bmatrix} \phi_1^p & \theta_1^p \\ 0 & 0 \end{bmatrix} \begin{bmatrix} S_{t-1} \\ \varepsilon_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \varepsilon_t, \quad \varepsilon_{t-1} \geq 0,\ \pi_{t-1} \geq d, \tag{16}$$

$$\begin{bmatrix} S_t \\ \varepsilon_t \end{bmatrix} = \begin{bmatrix} \phi_1^n & \theta_1^n \\ 0 & 0 \end{bmatrix} \begin{bmatrix} S_{t-1} \\ \varepsilon_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \varepsilon_t, \quad \varepsilon_{t-1} < 0,\ \pi_{t-1} \geq d, \tag{17}$$

$$\begin{bmatrix} S_t \\ \varepsilon_t \end{bmatrix} = \begin{bmatrix} \phi_2^p & \theta_2^p \\ 0 & 0 \end{bmatrix} \begin{bmatrix} S_{t-1} \\ \varepsilon_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \varepsilon_t, \quad \varepsilon_{t-1} \geq 0,\ \pi_{t-1} < d, \tag{18}$$

$$\begin{bmatrix} S_t \\ \varepsilon_t \end{bmatrix} = \begin{bmatrix} \phi_2^n & \theta_2^n \\ 0 & 0 \end{bmatrix} \begin{bmatrix} S_{t-1} \\ \varepsilon_{t-1} \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} \varepsilon_t, \quad \varepsilon_{t-1} < 0,\ \pi_{t-1} < d, \tag{19}$$

where

$$Q = E\left[\begin{bmatrix} \varepsilon_t \\ \varepsilon_t \end{bmatrix} [\,\varepsilon_t \quad \varepsilon_t\,] \right] = \sigma^2 \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}. \tag{20}$$

Let $x_t = [S_t\ \varepsilon_t]'$ and denote by Φ_1^p, Φ_1^n, Φ_2^p, and Φ_2^n the 2×2 matrices appearing in transition equations (16) to (19). The forecast value of x_t on the basis of information available through date $t-1$, denoted $x_{t|t-1}$, is given by

$$x_{t|t-1} = \Phi_1^p x_{t-1|t-1}, \quad x_{t-1|t-1}^{[2]} \geq 0,\ \pi_{t-1} \geq d, \tag{21}$$

$$x_{t|t-1} = \Phi_1^n x_{t-1|t-1}, \quad x_{t-1|t-1}^{[2]} < 0, \ \pi_{t-1} \geq d, \tag{22}$$

$$x_{t|t-1} = \Phi_2^p x_{t-1|t-1}, \quad x_{t-1|t-1}^{[2]} \geq 0, \ \pi_{t-1} < d, \tag{23}$$

$$x_{t|t-1} = \Phi_2^n x_{t-1|t-1}, \quad x_{t-1|t-1}^{[2]} < 0, \ \pi_{t-1} < d, \tag{24}$$

where $x_{t-1|t-1}^{[2]}$ corresponds to the second element of $x_{t-1|t-1}$. Let z_t denote the observations on output growth, Δy_τ, and inflation, π_τ, up to date t. Then the distribution of x_t conditional on z_{t-1} is normal, with mean $x_{t|t-1}$ and variance $P_{t|t-1}$. The forecast equations for the conditional variance are given by

$$P_{t|t-1} = \Phi_1^p P_{t-1|t-1} \Phi_1^{p(t)} + Q, \quad x_{t-1|t-1}^{[2]} \geq 0, \ \pi_{t-1} \geq d, \tag{25}$$

$$P_{t|t-1} = \Phi_1^n P_{t-1|t-1} \Phi_1^{n(t)} + Q, \quad x_{t-1|t-1}^{[2]} < 0, \ \pi_{t-1} \geq d, \tag{26}$$

$$P_{t|t-1} = \Phi_2^p P_{t-1|t-1} \Phi_2^{p(t)} + Q, \quad x_{t-1|t-1}^{[2]} \geq 0, \ \pi_{t-1} < d, \tag{27}$$

$$P_{t|t-1} = \Phi_2^n P_{t-1|t-1} \Phi_2^{n(t)} + Q, \quad x_{t-1|t-1}^{[2]} < 0, \ \pi_{t-1} < d, \tag{28}$$

where the superscript (t) denotes the transpose matrix. The filter recursions are such that the estimates of $x_{t-1|t-1}$ and $P_{t-1|t-1}$ are computed before $x_{t|t-1}$ and $P_{t|t-1}$. Therefore, the choice of the appropriate Φ_j^i during each recursion of the filter is unambiguous.

The updating equations for x_t and P_t are given by

$$x_{t|t} = x_{t|t-1} + P_{t|t-1} H e_t / v_t, \tag{29}$$

$$P_{t|t} = P_{t|t-1} - P_{t|t-1} H H' P_{t|t-1} / v_t, \tag{30}$$

with

$$e_t = \Delta y_t - \mu - H' x_{t|t-1}, \tag{31}$$

and

$$v_t = H' P_{t|t-1} H, \tag{32}$$

where $H' = [1 \ 0]$.

Collect the $n = 11$ model parameters in the vector

$$\theta = (\mu, \sigma^2, \phi_1^p, \phi_1^n, \phi_2^p, \phi_2^n, \theta_1^p, \theta_1^n, \theta_2^p, \theta_2^n, d). \tag{33}$$

Under the assumed model, we have

$$\Delta y_t | z_{t-1}; \theta \sim N\left(\mu_t(\theta), \sigma_t^2(\theta)\right) \tag{34}$$

where

$$\mu_t(\theta) = \mu + H' P_{t|t-1} \tag{35}$$

$$\sigma_t^2(\theta) = v_t. \tag{36}$$

Given θ, the above equations can be iterated to compute the value of the log-likelihood function,

$$\sum_{t=1}^{T} \log f(\Delta y_t | z_{t-1}; \theta) = \tag{37}$$

$$-(Tn/2)\log(2\pi) - (1/2)\sum_{t=1}^{T} \sigma_t^2(\theta) - (1/2)\sum_{t=1}^{T} \left(\frac{\Delta y_t - \mu_t(\theta)}{\sigma_t}\right)^2, \tag{38}$$

which, in turn, can be numerically maximized to obtain quasi-maximum-likelihood parameter estimates.[6] See Hamilton (1994) for a general discussion of state-space models and Elwood (1998) for more details on the modified Kalman filter.

[6] Under the maintained assumption $\varepsilon_t \sim i.i.d.\ N(0, \sigma^2)$, these estimates correspond to the maximum-likelihood parameter estimates.

Bank of Canada Working Papers
Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. ***Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue.***

2001

2001-21	A Consistent Bootstrap Test for Conditional Density Functions with Time-Dependent Data	F. Li and G. Tkacz
2001-20	The Resolution of International Financial Crises: Private Finance and Public Funds	A. Haldane and M. Kruger
2001-19	Employment Effects of Restructuring in the Public Sector in North America	P. Fenton, I. Ip, and G. Wright
2001-18	Evaluating Factor Models: An Application to Forecasting Inflation in Canada	M.-A. Gosselin and G. Tkacz
2001-17	Why Do Central Banks Smooth Interest Rates?	G. Srour
2001-16	Implications of Uncertainty about Long-Run Inflation and the Price Level	G. Stuber
2001-15	Affine Term-Structure Models: Theory and Implementation	D.J. Bolder
2001-14	L'effet de la richesse sur la consommation aux États-Unis	Y. Desnoyers
2001-13	Predetermined Prices and the Persistent Effects of Money on Output	M.B. Devereux and J. Yetman
2001-12	Evaluating Linear and Non-Linear Time-Varying Forecast-Combination Methods	F. Li and G. Tkacz
2001-11	Gaining Credibility for Inflation Targets	J. Yetman
2001-10	The Future Prospects for National Financial Markets and Trading Centres	C. Gaa, S. Lumpkin, R. Ogrodnik, and P. Thurlow
2001-9	Testing for a Structural Break in the Volatility of Real GDP Growth in Canada	A. Debs
2001-8	How Rigid Are Nominal-Wage Rates?	A. Crawford
2001-7	Downward Nominal-Wage Rigidity: Micro Evidence from Tobit Models	A. Crawford and G. Wright
2001-6	The Zero Bound on Nominal Interest Rates: How Important Is It?	D. Amirault and B. O'Reilly

Copies and a complete list of working papers are available from:
Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: http://www.bankofcanada.ca

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : http://www.banqueducanada.ca